arXiv:1108.6032v1 [math.ST] 30 Aug 2011

Likelihood inference for Archimedean copulas

Marius Hofert[1], Martin Mächler[2], Alexander J. McNeil[3]

2011-08-31

Abstract

Explicit functional forms for the generator derivatives of well-known one-parameter Archimedean copulas are derived. These derivatives are essential for likelihood inference as they appear in the copula density, conditional distribution functions, or the Kendall distribution function. They are also required for several asymmetric extensions of Archimedean copulas such as Khoudraji-transformed Archimedean copulas. Access to the generator derivatives makes maximum-likelihood estimation for Archimedean copulas feasible in terms of both precision and run time, even in large dimensions. It is shown by simulation that the root mean squared error is decreasing in the dimension. This decrease is of the same order as the decrease in sample size. Furthermore, confidence intervals for the parameter vector are derived. Moreover, extensions to multi-parameter Archimedean families are given. All presented methods are implemented in the open-source R package `nacopula` and can thus easily be accessed and studied.

Keywords

Archimedean copulas, maximum-likelihood estimation, confidence intervals, multi-parameter families.

MSC2010

62H12, 62F10, 62H99, 65C60.

Submitted

1 Introduction

The well-known class of *Archimedean copulas* consists of copulas of the form

$$C(\boldsymbol{u}) = \psi(\psi(u_1) + \cdots + \psi(u_d)),\ \boldsymbol{u} \in [0,1]^d,$$

with *generator* ψ. In practical applications, ψ belongs to a parametric family $(\psi_{\boldsymbol{\theta}})_{\boldsymbol{\theta}\in\Theta}$ whose parameter vector $\boldsymbol{\theta}$ needs to be estimated.

There are several known approaches for estimating parametric Archimedean copula families; see Hofert et al. (2011) for an overview and a comparison of some estimators.

[1]RiskLab, Department of Mathematics, ETH Zurich, 8092 Zurich, Switzerland, marius.hofert@math.ethz.ch. The author (Willis Research Fellow) thanks Willis Re for financial support while this work was being completed.
[2]Seminar für Statistik, ETH Zurich, 8092 Zurich, Switzerland, maechler@stat.math.ethz.ch
[3]Department of Actuarial Mathematics and Statistics, Heriot-Watt University, Edinburgh, EH14 4AS, Scotland, A.J.McNeil@hw.ac.uk

In the work at hand, we consider a (semi-)parametric estimation approach based on the likelihood. There are two significant obstacles to overcome. The first one is to derive tractable algebraic expressions for the generator derivatives and thus the copula density. The second is to evaluate these expressions efficiently in terms of both precision and run time.

Although the density of an Archimedean copula has an explicit form in theory, accessing the required derivatives is known to be challenging, especially in large dimensions. For example, Berg and Aas (2009) mention that for Archimedean copulas it is not straightforward to derive the density in general for all parametric families. For the Gumbel family, they say that one has to resort to a computer algebra system, such as Mathematica or the function `D` in R, to derive the d-dimensional density. Note that computations based on computer algebra systems often fail already in low dimensions. Even if a theoretical formula can be computed, the numerical evaluation of such (typically lengthy) formulas is prone to errors since they are not given in a numerically tractable form. This often requires to work with a large number of significant digits which is typically far too slow to be applied in large-scale simulation studies (for example, to access the quality of goodness-of-fit testing procedures). Furthermore, as we will point out below, results obtained by computer algebra systems can be unreliable.

Generator derivatives for some important Archimedean families can be found in Shi (1995), Barbe et al. (1996), and Wu et al. (2007), however, in recursive form. In this work, we derive explicit formulas for the generator derivatives of well-known Archimedean families in any dimension. These derivatives are interesting in their own right, for example, for accessing densities, for building conditional distribution functions, or for evaluating the Kendall distribution function. They can also be used to explicitly compute densities of asymmetric extensions of Archimedean copulas such as Khoudraji-transformed Archimedean copulas.

We then tackle the problem of maximum-likelihood estimation for Archimedean copulas for these families. Focus is put on large, say ten to one hundred, dimensions since they are the most relevant in practice; see Embrechts and Hofert (2011). Note that the considered Gumbel family is also an extreme value copula, for which densities in general are rarely known. Hofert et al. (2011) show the excellent performance of the maximum-likelihood estimator as measured by both precision and run time in a large-scale comparison with various other estimators up to dimension one hundred. Furthermore, to add transparency, all the algorithms used in this paper are implemented in the open source R package `nacopula`, so that the interested reader can study the non-trivial details of the numerical implementation and the numerous tests conducted in more detail. In the work at hand, we also consider examples of multi-parameter Archimedean families. In contrast to method-of-moments-like estimation procedures such as the one based on Kendall's tau, maximum-likelihood estimation is not limited to the one-parameter case. Furthermore, we address the problem of computing initial intervals for the optimization of the log-likelihood for the multi-parameter Archimedean families considered. Additionally, we show how confidence intervals for the copula parameter vector can be constructed.

The paper is organized as follows. In Section 2, we briefly recall the notion of Archimedean copulas and the families considered. Section 3 presents explicit functional

Submitted

forms of the generator derivatives of these families and the corresponding copula densities are derived. In Section 4, the root mean squared error is investigated as a function of the dimension. Section 5 presents methods for constructing confidence intervals for the copula parameter vector. In Section 6 we address extensions to multi-parameter Archimedean families, including a strategy for computing initial intervals and two examples of two-parameter families. Finally, Section 7 concludes.

2 Archimedean copulas

Definition 2.1
An *(Archimedean) generator* is a continuous, decreasing function $\psi : [0, \infty] \to [0, 1]$ which satisfies $\psi(0) = 1$, $\psi(\infty) = \lim_{t\to\infty} \psi(t) = 0$, and which is strictly decreasing on $[0, \inf\{t : \psi(t) = 0\}]$. A d-dimensional copula C is called *Archimedean* if it permits the representation

$$C(\boldsymbol{u}) = \psi(\psi^{-1}(u_1) + \cdots + \psi^{-1}(u_d)), \ \boldsymbol{u} \in [0, 1]^d, \tag{1}$$

for some generator ψ with inverse $\psi^{-1} : [0, 1] \to [0, \infty]$, where $\psi^{-1}(0) = \inf\{t : \psi(t) = 0\}$.

McNeil and Nešlehová (2009) show that a generator defines an Archimedean copula if and only if ψ is *d-monotone*, meaning that ψ is continuous on $[0, \infty]$, admits derivatives up to the order $d - 2$ satisfying $(-1)^k \frac{d^k}{dt^k}\psi(t) \geq 0$ for all $k \in \{0, \ldots, d - 2\}$, $t \in (0, \infty)$, and $(-1)^{d-2} \frac{d^{d-2}}{dt^{d-2}}\psi(t)$ is decreasing and convex on $(0, \infty)$.

According to McNeil and Nešlehová (2009), an Archimedean copula C admits a density c if and only if $\psi^{(d-1)}$ exists and is absolutely continuous on $(0, \infty)$. In this case, c is given by

$$c(\boldsymbol{u}) = \psi^{(d)}(t(\boldsymbol{u})) \prod_{j=1}^{d} (\psi^{-1})'(u_j), \ \boldsymbol{u} \in (0, 1)^d, \tag{2}$$

where $t(\boldsymbol{u}) = \sum_{j=1}^{d} \psi(u_j)$.

We mainly assume ψ to be *completely monotone*, meaning that ψ is continuous on $[0, \infty]$ and $(-1)^k \frac{d^k}{dt^k}\psi(t) \geq 0$ for all $k \in \mathbb{N}_0$, $t \in (0, \infty)$, so that ψ is the Laplace-Stieltjes transform of a distribution function F on the positive real line, that is, $\psi = \mathcal{LS}[F]$; see Bernstein's Theorem in Feller (1971, p. 439). The class of all such generators is denoted by Ψ_∞ and it is clear that a $\psi \in \Psi_\infty$ generates an Archimedean copula in any dimensions d and that its density exists.

There are several well-known parametric generator families; see Nelsen (2007, pp. 116), also referred to as *Archimedean families*. Among the most widely used in applications are those of Ali-Mikhail-Haq ("A"), Clayton ("C"), Frank ("F"), Gumbel ("G"), and Joe ("J"); see Table 1. We consider these families as working examples throughout this work. Detailed information about the corresponding distribution functions F is given in Hofert (2011b) and references therein. Note that these one-parameter families can be extended to allow for more parameters, for example, via outer power transformations.

Submitted

Family	Parameter	$\psi(t)$	$V \sim F = \mathcal{LS}^{-1}[\psi]$
A	$\theta \in [0, 1)$	$(1-\theta)/(\exp(t)-\theta)$	$\mathrm{Geo}(1-\theta)$
C	$\theta \in (0, \infty)$	$(1+t)^{-1/\theta}$	$\Gamma(1/\theta, 1)$
F	$\theta \in (0, \infty)$	$-\log(1-(1-e^{-\theta})\exp(-t))/\theta$	$\mathrm{Log}(1-e^{-\theta})$
G	$\theta \in [1, \infty)$	$\exp(-t^{1/\theta})$	$\mathrm{S}(1/\theta, 1, \cos^{\theta}(\pi/(2\theta)), \mathbb{1}_{\{\theta=1\}}; 1)$
J	$\theta \in [1, \infty)$	$1-(1-\exp(-t))^{1/\theta}$	$\mathrm{Sibuya}(1/\theta)$

Table 1 Well-known one-parameter Archimedean generators ψ with corresponding distributions $F = \mathcal{LS}^{-1}[\psi]$.

Furthermore, there are Archimedean families which are naturally given by more than a single parameter. Examples for both cases are given in Section 6.

Table 2 summarizes properties concerning Kendall's tau and the tail-dependence coefficients; see Joe (1997, p. 91), Joe and Hu (1996), and Nelsen (2007, p. 214) for the investigated Archimedean families. Here, $D_1(\theta) = \int_0^\theta t/(\exp(t)-1)\,dt/\theta$ denotes the *Debye function of order one*. Note that these properties are often of interest in order to choose a suitable model which is then estimated. The construction of initial intervals in Section 6.1 for the optimization of the likelihood is based on Kendall's tau.

Family	τ	λ_L	λ_U
A	$1-2(\theta+(1-\theta)^2\log(1-\theta))/(3\theta^2)$	0	0
C	$\theta/(\theta+2)$	$2^{-1/\theta}$	0
F	$1+4(D_1(\theta)-1)/\theta$	0	0
G	$(\theta-1)/\theta$	0	$2-2^{1/\theta}$
J	$1-4\sum_{k=1}^{\infty} 1/(k(\theta k+2)(\theta(k-1)+2))$	0	$2-2^{1/\theta}$

Table 2 Kendall's tau and tail-dependence coefficients.

Submitted

3 Maximum-likelihood estimation for Archimedean copulas

3.1 The pseudo maximum-likelihood estimator

Assume that we have given realizations $\boldsymbol{x}_i$, $i \in \{1, \dots, n\}$, of independent and identically distributed ("i.i.d.") random vectors $\boldsymbol{X}_i$, $i \in \{1, \dots, n\}$, from a joint distribution function H with Archimedean copula C generated by ψ and corresponding density c. The generator ψ is assumed to belong to a parametric family $(\psi_{\boldsymbol{\theta}})_{\boldsymbol{\theta} \in \Theta}$ with parameter vector $\boldsymbol{\theta} \in \Theta \subseteq \mathbb{R}^p$, $p \in \mathbb{N}$, and the true but unknown vector is $\boldsymbol{\theta}_0$ (similarly, $C = C_{\boldsymbol{\theta}_0}$ and $c = c_{\boldsymbol{\theta}_0}$). As usual, random vectors or random variables are denoted by upper-case letters, their realizations by lower-case letters.

Before estimating $\boldsymbol{\theta}_0$, the first step is usually to estimate the marginal distribution functions. In a second step, one then estimates $\boldsymbol{\theta}_0$. This two-step approach is typically

much easier to accomplish than estimating the parameters of the marginal distribution functions and the copula parameter vector simultaneously. Estimating the marginal distribution functions can be done either parametrically or non-parametrically. Based on maximum-likelihood estimation, the former approach is suggested by Joe and Xu (1996) and is known as *inference functions for margins*. The latter approach is known as *pseudo maximum-likelihood estimation* and is suggested by Genest et al. (1995); see Kim et al. (2007) for a comparison of maximum-likelihood estimation, the method of inference functions for margins, and pseudo maximum-likelihood estimation.

Following pseudo maximum-likelihood estimation, the marginal distribution functions are estimated by their empirical distribution functions $\hat{F}_{nj}(x) = \frac{1}{n}\sum_{k=1}^{n} \mathbb{1}_{\{x_{kj} \le x\}}$, $j \in \{1, \dots, n\}$, leading to the so-called *pseudo-observations* $\hat{\boldsymbol{u}}_i = (\hat{u}_{i1}, \dots, \hat{u}_{id})^\mathsf{T}$, $i \in \{1, \dots, n\}$, where

$$\hat{u}_{ij} = \frac{n}{n+1}\hat{F}_{nj}(x_{ij}) = \frac{r_{ij}}{n+1}, \ i \in \{1, \dots, n\}, \ j \in \{1, \dots, d\}. \tag{3}$$

Here, for each $j \in \{1, \dots, d\}$, r_{ij} denotes the rank of x_{ij} among all x_{kj}, $k \in \{1, \dots, n\}$. The asymptotically negligible scaling factor of $n/(n+1)$ is used to force the variates to fall inside the open unit hypercube to avoid problems with density evaluation at the boundaries of $[0,1]^d$. As usual, the pseudo-observations are interpreted as realizations of a random sample from C (despite known issues of this interpretation such as the fact that the pseudo-observations are neither realizations of perfectly independent random vectors nor that the components are perfectly following a univariate standard uniform distribution) based on which the copula parameter vector $\boldsymbol{\theta}_0$ is estimated.

3.2 Likelihood theory

Maximum-likelihood estimation is based on the following theory. Given realizations $\boldsymbol{u}_i$, $i \in \{1, \dots, n\}$, of a random sample $\boldsymbol{U}_i$, $i \in \{1, \dots, n\}$, from the copula C (in practice, $\boldsymbol{u}_i$ is taken as $\hat{\boldsymbol{u}}_i$, $i \in \{1, \dots, n\}$, in (3)), the *likelihood* and *log-likelihood* are defined by

$$L(\boldsymbol{\theta}; \boldsymbol{u}_1, \dots, \boldsymbol{u}_n) = \prod_{i=1}^{n} c_{\boldsymbol{\theta}}(\boldsymbol{u}_i) \quad \text{and} \quad l(\boldsymbol{\theta}; \boldsymbol{u}_1, \dots, \boldsymbol{u}_n) = \sum_{i=1}^{n} l(\boldsymbol{\theta}; \boldsymbol{u}_i),$$

respectively, where

$$l(\boldsymbol{\theta}; \boldsymbol{u}_i) = \log c_{\boldsymbol{\theta}}(\boldsymbol{u}_i) = \log\big((-1)^d \psi_{\boldsymbol{\theta}}^{(d)}(t_{\boldsymbol{\theta}}(\boldsymbol{u}))\big) + \sum_{j=1}^{d} \log(-(\psi_{\boldsymbol{\theta}}^{-1})'(u_{ij})).$$

Here, the subscript $\boldsymbol{\theta}$ of $t(\boldsymbol{u})$ is used to stress the dependence of $t(\boldsymbol{u})$ on $\boldsymbol{\theta}$. The *maximum-likelihood estimator* $\hat{\boldsymbol{\theta}}_n = \hat{\boldsymbol{\theta}}_n(\boldsymbol{u}_1, \dots, \boldsymbol{u}_n)$ can thus be found by solving the optimization problem

$$\hat{\boldsymbol{\theta}}_n = \operatorname*{argsup}_{\boldsymbol{\theta} \in \Theta} l(\boldsymbol{\theta}; \boldsymbol{u}_1, \dots, \boldsymbol{u}_n).$$

This optimization is typically done numerically.

Submitted

Assuming the derivatives to exist, the *score function* is defined as

$$s_{\boldsymbol{\theta}}(\boldsymbol{u}) = \nabla l(\boldsymbol{\theta}; \boldsymbol{u}) = \left(\frac{\partial}{\partial \theta_1} l(\boldsymbol{\theta}; \boldsymbol{u}), \dots, \frac{\partial}{\partial \theta_p} l(\boldsymbol{\theta}; \boldsymbol{u}) \right)^\mathsf{T}$$

and the *Fisher information* is

$$I(\boldsymbol{\theta}) = \mathbb{E}_{\boldsymbol{\theta}}[s_{\boldsymbol{\theta}}(\boldsymbol{U}) s_{\boldsymbol{\theta}}(\boldsymbol{U})^\mathsf{T}] = \mathbb{E}_{\boldsymbol{\theta}}\left[\left(\frac{\partial}{\partial \theta_i} l(\boldsymbol{\theta}; \boldsymbol{u}) \frac{\partial}{\partial \theta_j} l(\boldsymbol{\theta}; \boldsymbol{u}) \right)_{i,j \in \{1,\dots,p\}} \right]$$

for $\boldsymbol{U} \sim C$.

Under regularity conditions (see Cox and Hinkley (1974, p. 281), Rohatgi (1976, pp. 384), Serfling (1980, pp. 144), Newey and McFadden (1994, p. 2146), Schervish (1995, p. 421), Lehmann and Casella (1998, p. 449), van der Vaart (2000, pp. 51), Bickel and Doksum (2000, p. 386), or Davison (2003, p. 118)), the following result holds.

Theorem 3.1

(1) (Strong) consistency of maximum-likelihood estimators:

$$\hat{\boldsymbol{\theta}}_n = \hat{\boldsymbol{\theta}}_n(\boldsymbol{U}_1, \dots, \boldsymbol{U}_n) \xrightarrow[\text{a.s.}]{P} \boldsymbol{\theta}_0 \ (n \to \infty).$$

(2) Asymptotic normality of maximum-likelihood estimators:

$$\sqrt{n}\, I(\boldsymbol{\theta}_0)^{1/2} (\hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}_0) \xrightarrow{d} N(\boldsymbol{0}, I_p),$$

where I_p denotes the identity matrix in $\mathbb{R}^{p \times p}$.

3.3 Generator derivatives and copula density

Applying maximum-likelihood estimation requires an efficient strategy for evaluating the (log-)density of the parametric Archimedean copula family to be estimated. The most important part is to know how to access the generator derivatives. As mentioned in the introduction, this requires to know both a tractable algebraic form of the derivatives and a procedure to numerically evaluate the formulas in an efficient way in terms of precision and run time.

As mentioned in the introduction, it is often stated that a computer algebra system can be used to access a generator's derivatives. Such an approach has typically two major flaws:

(1) It is not trivial and sometimes not possible for a computer algebra system to find derivatives of higher order;

(2) Even if formulas are obtained, they are usually not provided in a form which is both numerically stable and sufficiently fast to evaluate.

We experienced these flaws when we tried to access the 50th derivative of a Gumbel generator $\psi_\theta(t)$ with parameter $\theta = 1.25$ at $t = 15$. On a MacBook Pro running Max OS

Submitted

X 10.6.6, we aborted Mathematica 8 after ten minutes without obtaining a result. Maple 14 lead to the values 10 628, -29 800, and others (without warning) when computing $\psi_{1.25}^{(50)}(15)$ several times. Note the chaotic behavior of this deterministic problem; the values should of course be equal and positive! MATLAB 7.11.0 did return the correct value of (roughly) 1057, but failed to access $\psi_{1.25}^{(100)}(15)$ (aborted after ten minutes). Let us stress that carelessly using such programs in simulations may lead to wrong results. Apart from numerical issues, the formulas for the derivatives obtained from computer algebra systems can become quite large and thus rather slow to evaluate. They are therefore not suitable in large-scale simulation studies, for example, for goodness-of-fit tests (or simulations of their performance) involving a parametric bootstrap.

In the following theorem we derive explicit formulas for the generator derivatives for all Archimedean families given in Table 1.

Theorem 3.2

(1) For the family of Ali-Mikhail-Haq,

$$(-1)^d \psi_\theta^{(d)}(t) = \frac{1-\theta}{\theta} \operatorname{Li}_{-d}(\theta \exp(-t)),\ t \in (0,\infty),\ d \in \mathbb{N}_0,$$

where $\operatorname{Li}_s(z)$ denotes the *polylogarithm of order s at z*.

(2) For the family of Clayton,

$$(-1)^d \psi_\theta^{(d)}(t) = (d-1+1/\theta)_d (1+t)^{-(d+1/\theta)},\ t \in (0,\infty),\ d \in \mathbb{N}_0,$$

where $(d-1+1/\theta)_d = \prod_{k=0}^{d-1}(k+1/\theta) = \frac{\Gamma(d+1/\theta)}{\Gamma(1/\theta)}$ denotes the falling factorial.

(3) For the family of Frank,

$$(-1)^d \psi_\theta^{(d)}(t) = \frac{1}{\theta} \operatorname{Li}_{-(d-1)}((1-e^{-\theta}) \exp(-t)),\ t \in (0,\infty),\ d \in \mathbb{N}_0.$$

(4) For the family of Gumbel,

$$(-1)^d \psi_\theta^{(d)}(t) = \frac{\psi_\theta(t)}{t^d} P_{d,\theta}^{\mathrm{G}}(t^\alpha),\ t \in (0,\infty),\ d \in \mathbb{N},$$

where

$$P_{d,\theta}^{\mathrm{G}}(x) = \sum_{k=1}^{d} a_{dk}^{\mathrm{G}}(\theta) x^k,$$

$$a_{dk}^{\mathrm{G}}(\theta) = (-1)^{d-k} \sum_{j=k}^{d} \theta^{-j} s(d,j) S(j,k) = \frac{d!}{k!} \sum_{j=1}^{k} \binom{k}{j} \binom{\alpha j}{d} (-1)^{d-j},\ k \in \{1,\dots,d\},$$

and s and S denote the *Stirling numbers of the first kind* and the *second kind*, respectively.

Submitted

(5) For the family of Joe,

$$(-1)^d \psi_\theta^{(d)}(t) = \frac{\exp(-t)}{\theta(1-\exp(-t))^{1-1/\theta}} P_{d,\theta}^{\mathrm{J}}\Big(\frac{\exp(-t)}{1-\exp(-t)}\Big),\ t \in (0,\infty),\ d \in \mathbb{N},$$

where

$$P_{d,\theta}^{\mathrm{J}}(x) = \sum_{k=1}^{d} a_{dk}^{\mathrm{J}}(\theta) x^{k-1},$$
$$a_{dk}^{\mathrm{J}}(\theta) = S(d,k)(k-1-1/\theta)_{k-1} = S(d,k)\frac{\Gamma(k-\alpha)}{\Gamma(1-\alpha)},\ k \in \{1,\dots,d\}.$$

Proof

(1) The generator of the Archimedean family of Ali-Mikhail-Haq is of the form $\psi_\theta(t) = \sum_{k=1}^{\infty} p_k \exp(-kt)$, $t \in [0,\infty)$, with probability mass function $(p_k)_{k=1}^{\infty}$ as given in Table 1. This implies that $(-1)^d \psi_\theta^{(d)}(t) = \sum_{k=1}^{\infty} p_k k^d \exp(-kt)$ from which the statement easily follows from the definition of the polylogarithm as $\mathrm{Li}_s(z) = \sum_{k=1}^{\infty} z^k/k^s$.

(2) The result for Clayton is straightforward to obtain by taking the derivatives.

(3) Similar to (1).

(4) Now consider Gumbel's family. Writing the generator in terms of the exponential series and differentiating the summands, leads to $\psi_\theta^{(d)}(t) = \sum_{k=1}^{\infty}(-1)^k/k!(\alpha k)_d t^{\alpha k - d}$, where $\alpha = 1/\theta$. Since for $d \in \mathbb{N}$, $(\alpha k)_d = \sum_{j=1}^{d} s(d,j)(\alpha k)^j$, one obtains $\psi_\theta^{(d)}(t) = t^{-d}\sum_{k=1}^{\infty}(-t^\alpha)^k/k! \sum_{j=1}^{d} s(d,j)(\alpha k)^j = t^{-d}\sum_{j=1}^{d}\alpha^j s(d,j)\sum_{k=1}^{\infty} k^j(-t^\alpha)^k/k!$. Note that $\exp(-x)\sum_{k=0}^{\infty} k^j x^k/k!$ is the *jth exponential polynomial* and equals $\sum_{k=0}^{j} S(j,k)\cdot x^k$; see Boyadzhiev (2009). With $x = -t^\alpha$ and noting that the summand for $k=0$ is zero, we obtain $\psi_\theta^{(d)}(t) = \psi_\theta(t) t^{-d}\sum_{j=1}^{d}\alpha^j s(d,j)\sum_{k=1}^{j} S(j,k)(-t^\alpha)^k$. Interchanging the order of summation leads to $\psi_\theta^{(d)}(t) = \psi_\theta(t) t^{-d}\sum_{k=1}^{d}(-t^\alpha)^k \sum_{j=k}^{d}\alpha^j s(d,j)S(j,k)$ $= \psi_\theta(t)\sum_{k=1}^{d} t^{\alpha k-d}(-1)^k \sum_{j=k}^{d}\alpha^j s(d,j)S(j,k)$ from which the result about $(-1)^d\psi_\theta^{(d)}$ directly follows. For the last equality in the statement about $a_{dk}^{\mathrm{G}}(\theta)$ note that $k!/d! a_{dk}^{\mathrm{G}}(\theta) = (-1)^{d-k}k!/d!\sum_{j=0}^{d}\alpha^j s(d,j)S(j,k) = (-1)^{d-k}/d!\sum_{j=0}^{d}\alpha^j s(d,j)\sum_{l=0}^{k}$ $\cdot\binom{k}{l}(-1)^{k-l}l^j = (-1)^{d-k}/d!\sum_{l=0}^{k}\binom{k}{l}(-1)^{k-l}\sum_{j=0}^{d}(\alpha l)^j s(d,j) = (-1)^d\sum_{l=0}^{k}\binom{k}{l}\binom{\alpha l}{d}$ $\cdot(-1)^l$ from which the result follows.

(5) For Joe's family, $(-1)^d\psi_\theta^{(d)}(t) = (-1)^{d+1}\frac{d^d}{dt^d}(1-\exp(-t))^\alpha$, $d \in \mathbb{N}$, where $\alpha = 1/\theta$. Letting $x = \exp(-t)$, this equals $-(x\frac{d}{dx})^d(1-x)^\alpha$. The operator $x\frac{d}{dx}$ is investigated in Boyadzhiev (2009). It follows from the results there that $(-1)^d\psi_\theta^{(d)}(t) = -\sum_{k=1}^{d} S(d,k)(-x)^k(\alpha)_k(1-x)^{\alpha-k} = -(1-x)^\alpha\sum_{k=1}^{d} S(d,k)(\alpha)_k(-x/(1-x))^k$. Thus, $(-1)^d\psi_\theta^{(d)}(t) = \alpha(1-x)^\alpha\sum_{k=1}^{d} S(d,k)(k-1-\alpha)_{k-1}(x/(1-x))^k$. Resubstituting leads to the result as stated.

□

Submitted

With the notation as in Theorem 3.2, we obtain the following representations for the densities of the Archimedean families of Ali-Mikhail-Haq, Clayton, Frank, Gumbel, and Joe.

Corollary 3.3

(1) For the family of Ali-Mikhail-Haq,

$$c_\theta(\boldsymbol{u}) = \frac{(1-\theta)^{d+1}}{\theta^2} \frac{h_\theta^{\mathrm{A}}(\boldsymbol{u})}{\prod_{j=1}^d u_j^2} \operatorname{Li}_{-d}(h_\theta^{\mathrm{A}}(\boldsymbol{u})),$$

where $h_\theta^{\mathrm{A}}(\boldsymbol{u}) = \theta \prod_{j=1}^d \frac{u_j}{1-\theta(1-u_j)}$.

(2) For the family of Clayton,

$$c_\theta(\boldsymbol{u}) = \prod_{k=0}^{d-1}(\theta k + 1)\Big(\prod_{j=1}^d u_j\Big)^{-(1+\theta)} (1 + t_\theta(\boldsymbol{u}))^{-(d+1/\theta)}.$$

(3) For the family of Frank,

$$c_\theta(\boldsymbol{u}) = \left(\frac{\theta}{1-e^{-\theta}}\right)^{d-1} \operatorname{Li}_{-(d-1)}(h_\theta^{\mathrm{F}}(\boldsymbol{u})) \frac{\exp(-\theta \sum_{j=1}^d u_j)}{h_\theta^{\mathrm{F}}(\boldsymbol{u})},$$

where $h_\theta^{\mathrm{F}}(\boldsymbol{u}) = (1-e^{-\theta})^{1-d} \prod_{j=1}^d (1 - \exp(-\theta u_j))$.

(4) For the family of Gumbel,

$$c_\theta(\boldsymbol{u}) = \theta^d C_\theta(\boldsymbol{u}) \frac{\prod_{j=1}^d (-\log u_j)^{\theta-1}}{t_\theta(\boldsymbol{u})^d \prod_{j=1}^d u_j} P_{d,\theta}^{\mathrm{G}}(t_\theta(\boldsymbol{u})^{1/\theta}).$$

(5) For the family of Joe,

$$c_\theta(\boldsymbol{u}) = \theta^{d-1} \frac{\prod_{j=1}^d (1-u_j)^{\theta-1}}{(1-h_\theta^{\mathrm{J}}(\boldsymbol{u}))^{1-1/\theta}} P_{d,\theta}^{\mathrm{J}}\Big(\frac{h_\theta^{\mathrm{J}}(\boldsymbol{u})}{1-h_\theta^{\mathrm{J}}(\boldsymbol{u})}\Big),$$

where $h_\theta^{\mathrm{J}}(\boldsymbol{u}) = \prod_{j=1}^d (1-(1-u_j)^\theta)$.

Proof

The proof is tedious but straightforward to obtain from Formula (2) and the results from Theorem 3.2. □

The following remarks stress the importance of Theorem 3.2 and Corollary 3.3.

Remark 3.4

(1) Recursive formulas for the generator derivatives for some Archimedean families were presented by Barbe et al. (1996) and Wu et al. (2007). In contrast, Theorem 3.2 provides explicit formulas. As seen from Corollary 3.3, this allows us to explicitly compute the densities of the corresponding well-known and widely used Archimedean

Submitted

families, even in large dimensions. Furthermore, it allows us to compute conditional distribution functions based on these families and important statistical quantities such as the Kendall distribution function, which is of interest, for example, in goodness-of-fit testing; see Genest et al. (2006), Genest et al. (2009), or Hering and Hofert (2011). Among others, note that extreme value copulas rarely have an explicit form of the density, the important Gumbel family can now be added to this list.

(2) The derivatives presented in Theorem 3.2 also play an important role in asymmetric extensions of Archimedean copulas. For example, consider a *Khoudraji-transformed Archimedean copula* C, given by

$$C(\boldsymbol{u}) = C_\psi(u_1^{\alpha_1}, \dots, u_d^{\alpha_d})\Pi(u_1^{1-\alpha_1}, \dots, u_d^{1-\alpha_d}),$$

where C_ψ denotes an Archimedean copula generated by ψ, Π denotes the independence copula, and $\alpha_j \in [0,1]$, $j \in \{1, \dots, d\}$, are parameters. Given the generator derivatives, the density of a Khoudraji-transformed Archimedean copula is given by

$$c(\boldsymbol{u}) = \sum_{J \subseteq \{1,\dots,d\}} \psi_V^{(|J|)}\Big(\sum_{j=1}^d \psi_V^{-1}(u_j^{\alpha_j})\Big) \prod_{j \in J} \alpha_j (\psi_V^{-1})'(u_j^{\alpha_j}) \prod_{j \notin J} (1-\alpha_j) u_j^{-\alpha_j}.$$

This makes maximum likelihood estimation for these copulas feasible; see Hofert and Vrins (2011) for an application.

(3) As pointed out by Hofert (2010b, pp. 117), new Archimedean copulas are often constructed with simple transformations of the generators addressed in Theorem 3.2. The results in Theorem 3.2 might therefore carry over to other Archimedean families. In fact, one example for such a transformation is the outer power transformation addressed in Section 6.

(4) For an Archimedean generator ψ with unknown derivatives but known $F = \mathcal{LS}^{-1}[\psi]$, Hofert et al. (2011) suggested to approximate $(-1)^d \psi^{(d)}$ via

$$(-1)^d \psi^{(d)}(t) \approx \frac{1}{m} \sum_{k=1}^m V_k^d \exp(-V_k t), \ t \in (0, \infty),$$

where $V_k \sim F$, $k \in \{1, \dots, m\}$, are realizations of i.i.d. random variables following $F = \mathcal{LS}^{-1}[\psi]$. In the conducted simulation study, this approximation turned out to be quite accurate. Furthermore, it is typically straightforward to implement. However, such a Monte Carlo approach is of course slower than having a direct formula for the generator derivatives at hand.

4 Sample size n vs dimension d

The results of Hofert et al. (2011) indicate that the root mean squared error ("RMSE") is decreasing in the dimension for all other parameters (Archimedean family, dependence level measured by Kendall's tau, and sample size) fixed. This may be intuitive for exchangeable copulas since the curse of dimensionality is circumvented by symmetry. In this

Submitted

section we briefly investigate how the RMSE decreases in the dimension. Figure 1 shows a clear picture. For fixed Archimedean family (Ali-Mikhail-Haq ("AMH"), Clayton, Frank, Gumbel, and Joe), dependence level measured by Kendall's tau ($\tau \in \{0.25, 0.5, 0.75\}$), and sample size ($n \in \{20, 50, 100, 200\}$), the RMSE (estimated based on $N = 500$ replications) is decreasing in the dimension ($d \in \{5, 10, 20, 50, 100\}$). As the log-log plot further reveals, the decrease of the RMSE in the dimension d is of the same order as in the sample size n, that is, the mean squared error ("MSE") satisfies

$$\text{MSE} \propto \frac{1}{nd}.$$

Although this behavior in the sample size n is well-known, the behavior in the dimension d is rather impressive since it contradicts the findings of Weiß (2010), for example. In the latter work, conclusions are drawn based on simulations only involving small dimensions. In small dimensions, however, numerical problems are often not (regarded) as severe as in larger dimensions. Sometimes, they are simply not solved correctly. However, according to our experience, we believe that the larger the dimension of interest is, the more involved numerical issues typically are. This will certainly become more important in the future as applications are often high-dimensional.

5 Constructing confidence intervals

In this section, we describe different ways of how to obtain confidence intervals for the copula parameter vector $\boldsymbol{\theta}_0$.

5.1 Fisher information

It follows from Theorem 3.1 (2) that

$$(\hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}_0)^\mathsf{T} n I(\boldsymbol{\theta}_0)(\hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}_0) \xrightarrow{d} \chi^2_p \ (n \to \infty).$$

This result remains valid if $I(\boldsymbol{\theta}_0)$ is replace by a consistent estimator $\widehat{I(\boldsymbol{\theta}_0)}$. Therefore, an asymptotic $1 - \alpha$ *confidence region* for $\boldsymbol{\theta}_0$ is given by

$$\Big\{\boldsymbol{\theta} \in \Theta : (\hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta})^\mathsf{T} n \widehat{I(\boldsymbol{\theta}_0)} (\hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}) \leq q_{\chi^2_p}(1-\alpha)\Big\},$$

where $q_{\chi^2_p}(1-\alpha)$ denotes the $(1-\alpha)$-quantile of the chi-square distribution with p degrees of freedom. In the one-parameter case, an asymptotic $1 - \alpha$ *confidence interval* for θ_0 is given by

$$\left[\hat{\theta}_n - \frac{z_{1-\alpha/2}}{\sqrt{n\widehat{I(\theta)}}}, \hat{\theta}_n + \frac{z_{1-\alpha/2}}{\sqrt{n\widehat{I(\theta)}}}\right],$$

where $z_{1-\alpha/2} = \Phi^{-1}(1 - \alpha/2)$ denotes the $(1 - \alpha/2)$-quantile of the standard normal distribution function.

Submitted

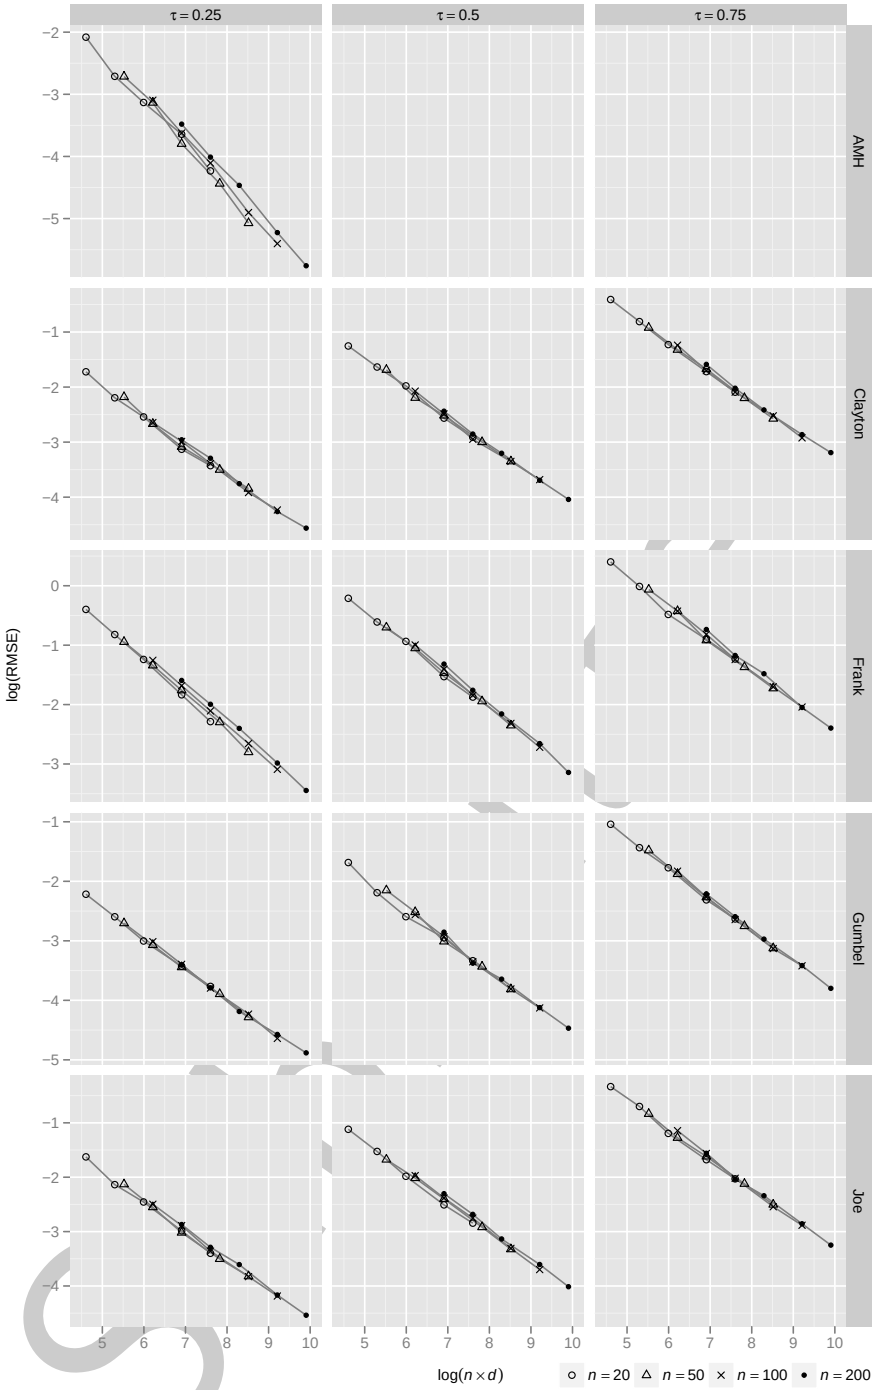

Figure 1 log-RMSE ($N = 500$ replications) as a function of the logarithm of $n \cdot d$. The plot indicates that the mean squared error satisfies MSE $\propto 1/(nd)$ for all families and dependencies. Note that the family of AMH is limited to $\tau \in [0, 1/3)$.

For the estimator $\widehat{I(\boldsymbol{\theta}_0)}$, there are several options, described in what follows. Assuming the derivatives to exist, the *observed information* is defined as

$$J(\boldsymbol{\theta}; \boldsymbol{u}_1, \ldots, \boldsymbol{u}_n) = -\nabla\nabla^\mathsf{T} l(\boldsymbol{\theta}; \boldsymbol{u}_1, \ldots, \boldsymbol{u}_n) = \sum_{i=1}^n -\nabla\nabla^\mathsf{T} l(\boldsymbol{\theta}; \boldsymbol{u}_i) \underset{p=1}{=} \sum_{i=1}^n -\frac{d^2}{d\theta^2} l(\theta; \boldsymbol{u}_i).$$

Under regularity conditions (see the references in Section 3.2), the Fisher information satisfies

$$I(\boldsymbol{\theta}) = \mathbb{E}[J(\boldsymbol{\theta}; \boldsymbol{U})] = \mathbb{E}[-\nabla\nabla^\mathsf{T} l(\boldsymbol{\theta}; \boldsymbol{U})] \underset{p=1}{=} \mathbb{E}\left[-\frac{d^2}{d\theta^2} l(\theta; \boldsymbol{U})\right],$$

that is, the Fisher information is the negative Hessian of the score function. From this and the definition of the Fisher information, the following choices for $\widehat{I(\boldsymbol{\theta}_0)}$ naturally arise (see also Newey and McFadden (1994, pp. 2157) including conditions for consistency):

$$I(\hat{\boldsymbol{\theta}}_n) = \mathbb{E}_{\hat{\boldsymbol{\theta}}_n}\left[s_{\hat{\boldsymbol{\theta}}_n}(\boldsymbol{U}) s_{\hat{\boldsymbol{\theta}}_n}(\boldsymbol{U})^\mathsf{T}\right] \tag{4}$$

$$\hat{I}^{(1)}(\hat{\boldsymbol{\theta}}_n) = \frac{1}{n}\sum_{i=1}^n s_{\hat{\boldsymbol{\theta}}_n}(\boldsymbol{u}_i) s_{\hat{\boldsymbol{\theta}}_n}(\boldsymbol{u}_i)^\mathsf{T} \tag{5}$$

$$\hat{I}^{(2)}(\hat{\boldsymbol{\theta}}_n) = \frac{1}{n}\sum_{i=1}^n J(\hat{\boldsymbol{\theta}}_n; \boldsymbol{u}_i) = \frac{1}{n}\sum_{i=1}^n -\nabla\nabla^\mathsf{T} l(\hat{\boldsymbol{\theta}}_n; \boldsymbol{u}_i) \tag{6}$$

The *expected information* $I(\hat{\boldsymbol{\theta}}_n)$ is often difficult to obtain. Furthermore, Efron and Hinkley (1978) argue for $\hat{I}^{(2)}(\hat{\boldsymbol{\theta}}_n)$ in favor of $I(\hat{\boldsymbol{\theta}}_n)$. The estimator $\hat{I}^{(1)}(\hat{\boldsymbol{\theta}}_n)$ is found much less in the literature, a reference being Newey and McFadden (1994, p. 2157). The reason why we state it here is that there are cases where the second-order partial derivatives are (much) more complicated to access than the first-order ones based on the score function.

The following proposition provides the score functions for the one-parameter Archimedean families given in Table 1.

Proposition 5.1

(1) For the family of Ali-Mikhail-Haq,

$$s_{\boldsymbol{\theta}}(\boldsymbol{u}) = -\frac{d+1}{1-\theta} - \frac{1}{\theta} + b_\theta^{\mathrm{A}}(\boldsymbol{u}) + \left(b_\theta^{\mathrm{A}}(\boldsymbol{u}) + \frac{1}{\theta}\right)\frac{\mathrm{Li}_{-(d+1)}(h_\theta^{\mathrm{A}}(\boldsymbol{u}))}{\mathrm{Li}_{-d}(h_\theta^{\mathrm{A}}(\boldsymbol{u}))},$$

where $b_\theta^{\mathrm{A}}(\boldsymbol{u}) = \sum_{j=1}^d \frac{1-u_j}{1-\theta(1-u_j)}$.

(2) For the family of Clayton,

$$s_{\boldsymbol{\theta}}(\boldsymbol{u}) = \sum_{k=0}^{d-1} \frac{k}{\theta k + 1} - \sum_{j=1}^d \log u_j + \frac{1}{\theta^2}\log(1 + t_\theta(\boldsymbol{u})) - (d + 1/\theta)\frac{t_\theta(\boldsymbol{u})}{1 + t_\theta(\boldsymbol{u})}.$$

Submitted

(3) For the family of Frank,

$$s_{\boldsymbol{\theta}}(\boldsymbol{u}) = \frac{d-1}{\theta} - \sum_{j=1}^{d} \frac{u_j}{1-\exp(-\theta u_j)} + \Big(\sum_{j=1}^{d} \frac{u_j \exp(-\theta u_j)}{1-\exp(-\theta u_j)} - \frac{(d-1)e^{-\theta}}{1-e^{-\theta}}\Big) \cdot \frac{\mathrm{Li}_{-d}(h_\theta^{\mathrm{F}}(\boldsymbol{u}))}{\mathrm{Li}_{-(d-1)}(h_\theta^{\mathrm{F}}(\boldsymbol{u}))}.$$

(4) For the family of Gumbel,

$$s_{\boldsymbol{\theta}}(\boldsymbol{u}) = \frac{d - \log C_\theta(\boldsymbol{u}) \log(-\log C_\theta(\boldsymbol{u}))}{\theta} - b_\theta^{\mathrm{G}}(\boldsymbol{u})\Big(d - \frac{\log C_\theta(\boldsymbol{u})}{\theta}\Big) + \sum_{j=1}^{d} \log(-\log u_j) + \frac{Q_{d,\theta,\boldsymbol{u}}^{\mathrm{G}}(t_\theta(\boldsymbol{u})^{1/\theta})}{\theta P_{d,\theta}^{\mathrm{G}}(t_\theta(\boldsymbol{u})^{1/\theta})},$$

where $b_\theta^{\mathrm{G}}(\boldsymbol{u}) = \sum_{j=1}^{d} \log(-\log u_j)\psi^{-1}(u_j)/t_\theta(\boldsymbol{u})$ and $Q_{d,\theta,\boldsymbol{u}}^{\mathrm{G}}(x) = \sum_{k=1}^{d} a_{dk}^{\mathrm{G}}(\theta, \boldsymbol{u})x^k$ with $a_{dk}^{\mathrm{G}}(\theta, \boldsymbol{u}) = k(b_\theta^{\mathrm{G}}(\boldsymbol{u}) - \frac{1}{\theta}\log t_\theta(\boldsymbol{u}))a_{dk}^{\mathrm{G}}(\theta) - (-1)^{d-k}\sum_{j=k}^{d} j s(d,j)S(j,k)\theta^{-j}$.

(5) For the family of Joe,

$$s_{\boldsymbol{\theta}}(\boldsymbol{u}) = \frac{d-1}{\theta} + \sum_{j=1}^{d} \log(1-u_j) - \frac{\log(1-h_\theta^{\mathrm{J}}(\boldsymbol{u}))}{\theta^2} + \frac{(1-\frac{1}{\theta})h_\theta^{\mathrm{J}}(\boldsymbol{u})}{1-h_\theta^{\mathrm{J}}(\boldsymbol{u})} b_\theta^{\mathrm{J}}(\boldsymbol{u}) + \frac{Q_{d,\theta,\boldsymbol{u}}^{\mathrm{J}}(h_\theta^{\mathrm{J}}(\boldsymbol{u})/(1-h_\theta^{\mathrm{J}}(\boldsymbol{u})))}{P_{d,\theta}^{\mathrm{J}}(h_\theta^{\mathrm{J}}(\boldsymbol{u})/(1-h_\theta^{\mathrm{J}}(\boldsymbol{u})))},$$

where $b_\theta^{\mathrm{J}}(\boldsymbol{u}) = \sum_{j=1}^{d} \frac{-\log(1-u_j)(1-u_j)^\theta}{1-(1-u_j)^\theta}$ and $Q_{d,\theta,\boldsymbol{u}}^{\mathrm{J}}(x) = \sum_{k=1}^{d} a_{dk}^{\mathrm{J}}(\theta, \boldsymbol{u})x^{k-1}$ with $a_{dk}^{\mathrm{J}}(\theta, \boldsymbol{u}) = a_{dk}^{\mathrm{J}}(\theta)(\frac{1}{\theta}\sum_{j=1}^{k-1}\frac{1}{\theta j-1} + (k-1)b_\theta^{\mathrm{J}}(\boldsymbol{u})/(1-h_\theta^{\mathrm{J}}(\boldsymbol{u})))$.

Proof

The proof is quite tedious but straightforward to obtain from Corollary 3.3. □

5.2 Likelihood-based confidence intervals

Confidence regions or confidence intervals can also be constructed solely based on the likelihood function (without requiring its derivatives). For this, the *likelihood ratio statistic* is used, defined as

$$W(\boldsymbol{\theta}; \boldsymbol{u}_1, \ldots, \boldsymbol{u}_n) = 2(l(\hat{\boldsymbol{\theta}}_n; \boldsymbol{u}_1, \ldots, \boldsymbol{u}_n) - l(\boldsymbol{\theta}; \boldsymbol{u}_1, \ldots, \boldsymbol{u}_n)),$$

As Davison (2003, p. 126) notes, the likelihood ratio statistic asymptotically follows a chi-square distribution, meaning that

$$W(\boldsymbol{\theta}_0; \boldsymbol{U}_1, \ldots, \boldsymbol{U}_n) \xrightarrow{d} \chi_p^2 \ (n \to \infty).$$

Submitted

Based on this result, an asymptotic $1-\alpha$ *confidence region* for $\boldsymbol{\theta}_0$ is given by

$$\{\boldsymbol{\theta} \in \Theta \,:\, l(\boldsymbol{\theta}; \boldsymbol{u}_1, \ldots, \boldsymbol{u}_n) \geq l(\hat{\boldsymbol{\theta}}_n; \boldsymbol{u}_1, \ldots, \boldsymbol{u}_n) - q_{\chi^2_p}(1-\alpha)/2\}. \tag{7}$$

If only a sub-vector $\boldsymbol{\theta}_0^{\mathrm{i}} \in \Theta_{p^{\mathrm{i}}} \subseteq \mathbb{R}^{p^{\mathrm{i}}}$ of components of $\boldsymbol{\theta}_0 = ({\boldsymbol{\theta}_0^{\mathrm{i}}}^\mathsf{T}, {\boldsymbol{\theta}_0^{\mathrm{n}}}^\mathsf{T})^\mathsf{T}$ are of interest ($\boldsymbol{\theta}_0^{\mathrm{i}}$ and $\boldsymbol{\theta}_0^{\mathrm{n}}$ are referred to as *parameters of interest* and *nuisance parameters*, respectively), an asymptotic confidence region for $\boldsymbol{\theta}_0^i$ follows from a similar argument to before, based on the *profile log-likelihood*

$$l_{p_i}(\boldsymbol{\theta}^{\mathrm{i}}; \boldsymbol{u}_1, \ldots, \boldsymbol{u}_n) = \sup_{\boldsymbol{\theta}^{\mathrm{n}}} l\left(\begin{pmatrix}\boldsymbol{\theta}^{\mathrm{i}} \\ \boldsymbol{\theta}^{\mathrm{n}}\end{pmatrix}; \boldsymbol{u}_1, \ldots, \boldsymbol{u}_n\right) = l\left(\begin{pmatrix}\boldsymbol{\theta}^{\mathrm{i}} \\ \hat{\boldsymbol{\theta}}_n^{\mathrm{n},\boldsymbol{\theta}^{\mathrm{i}}}\end{pmatrix}; \boldsymbol{u}_1, \ldots, \boldsymbol{u}_n\right),$$

where $\hat{\boldsymbol{\theta}}_n^{\mathrm{n},\boldsymbol{\theta}^{\mathrm{i}}}$ is the maximum-likelihood estimator of $\boldsymbol{\theta}_0^{\mathrm{n}}$ given $\boldsymbol{\theta}^{\mathrm{i}}$. Under regularity conditions, the *generalized likelihood ratio statistic*

$$W_{p^{\mathrm{i}}}(\boldsymbol{\theta}^{\mathrm{i}}; \boldsymbol{u}_1, \ldots, \boldsymbol{u}_n) = 2\left(l(\hat{\boldsymbol{\theta}}_n; \boldsymbol{u}_1, \ldots, \boldsymbol{u}_n) - l\left(\begin{pmatrix}\boldsymbol{\theta}^{\mathrm{i}} \\ \hat{\boldsymbol{\theta}}_n^{\mathrm{n},\boldsymbol{\theta}^{\mathrm{i}}}\end{pmatrix}; \boldsymbol{u}_1, \ldots, \boldsymbol{u}_n\right)\right)$$

satisfies

$$W_{p^{\mathrm{i}}}(\boldsymbol{\theta}_0^{\mathrm{i}}; \boldsymbol{U}_1, \ldots, \boldsymbol{U}_n) \xrightarrow{d} \chi^2_{p^{\mathrm{i}}} \; (n \to \infty).$$

An asymptotic $1-\alpha$ *confidence region* for $\boldsymbol{\theta}_0^{\mathrm{i}}$ is thus given by

$$\{\boldsymbol{\theta}^{\mathrm{i}} \in \Theta_{p^{\mathrm{i}}} \,:\, l_{p^{\mathrm{i}}}(\boldsymbol{\theta}^{\mathrm{i}}; \boldsymbol{u}_1, \ldots, \boldsymbol{u}_n) \geq l_{p^{\mathrm{i}}}(\hat{\boldsymbol{\theta}}_n^{\mathrm{i}}; \boldsymbol{u}_1, \ldots, \boldsymbol{u}_n) - q_{\chi^2_{p^{\mathrm{i}}}}(1-\alpha)/2\},$$

where

$$\hat{\boldsymbol{\theta}}_n^{\mathrm{i}} = \operatorname*{argsup}_{\boldsymbol{\theta}^{\mathrm{i}} \in \Theta^{\mathrm{i}}} l_{p^{\mathrm{i}}}(\boldsymbol{\theta}^{\mathrm{i}}; \boldsymbol{u}_1, \ldots, \boldsymbol{u}_n).$$

This will be used in Section 6 to construct confidence intervals for multi-parameter families.

Example 5.2

The left-hand side of Figure 2 shows the log-likelihood of a Clayton copula based on a 100-dimensional sample of size $n = 100$ with parameter $\theta_0 = 2$ such that the corresponding bivariate population version of Kendall's tau equals $\tau(\theta_0) = 0.5$. The maximum-likelihood estimator is denoted by $\hat{\theta}_n$ and the lower and upper endpoints of the likelihood-based 0.95 confidence interval by $\theta_l^{0.95}$ and $\theta_u^{0.95}$, respectively. The right-hand side of Figure 2 shows the profile likelihood plot for the same sample. Similarly for Figure 3 which shows the log-likelihood and profile likelihood plot for the 100-dimensional Gumbel family with parameter $\theta_0 = 2$ such that Kendall's tau equals $\tau(\theta_0) = 0.5$.

Submitted

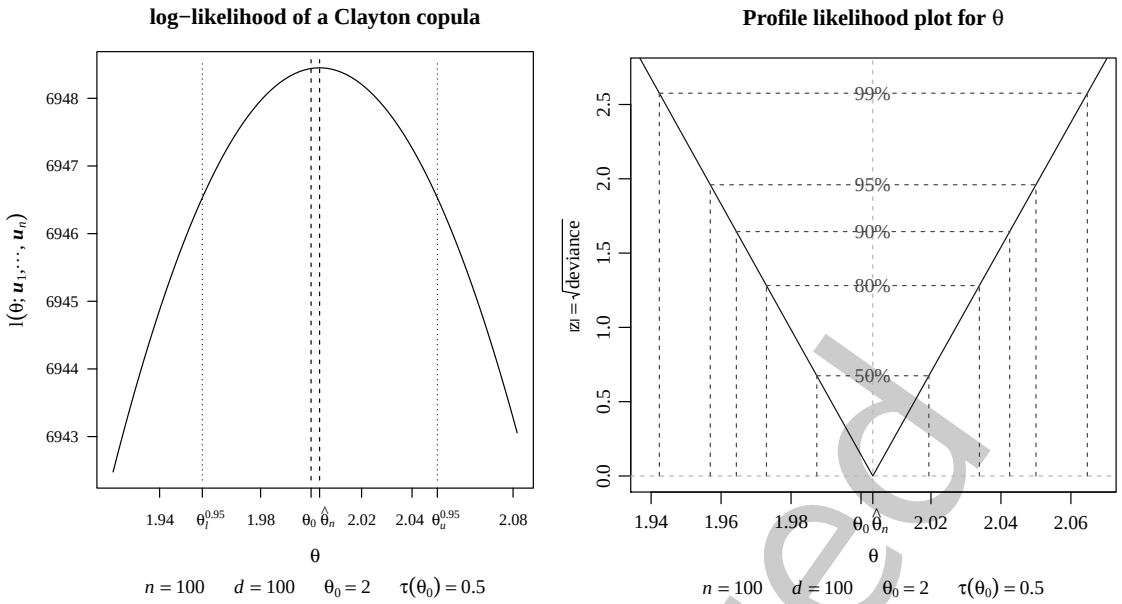

Figure 2 Plot of the log-likelihood of a Clayton copula (left) based on a sample of size $n = 100$ in dimension $d = 100$ with parameter $\theta_0 = 2$ such that Kendall's tau equals 0.5. Corresponding profile likelihood plot (right).

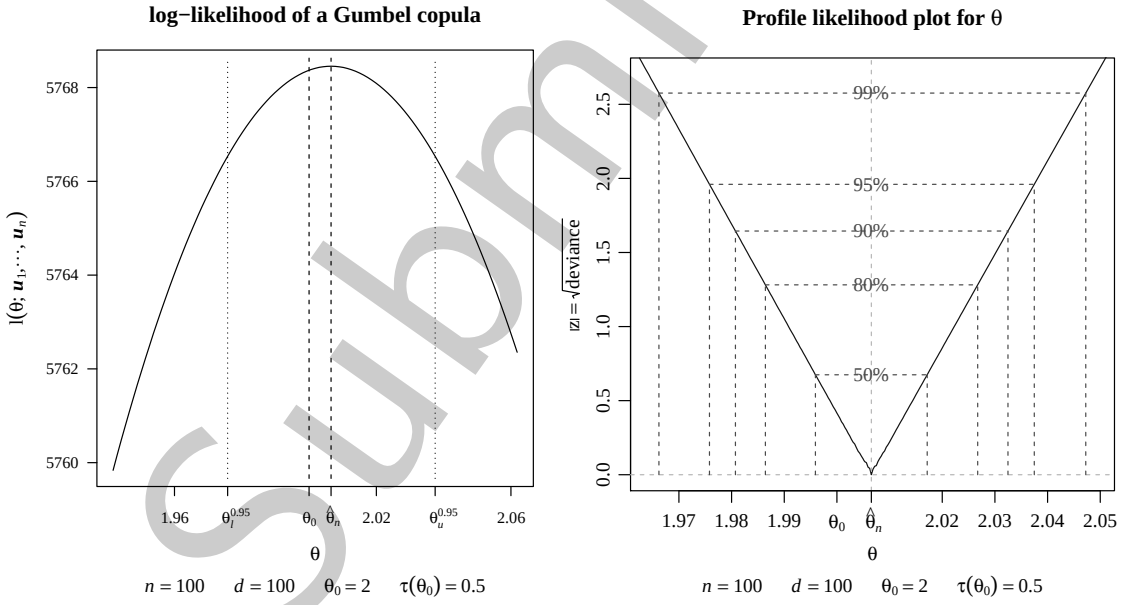

Figure 3 Plot of the log-likelihood of a Gumbel copula (left) based on a sample of size $n = 100$ in dimension $d = 100$ with parameter $\theta_0 = 2$ such that Kendall's tau equals 0.5. Corresponding profile likelihood plot (right).

Submitted

5.3 A simulation study to access the coverage probability

In this section, we compare the different approaches for obtaining (asymptotic) confidence regions and intervals. For this, we conduct a simulation study to access the coverage probability. The methods for obtaining confidence intervals based on the Fisher information are denoted by "$I(\hat{\theta}_n)$" for (4), "$\hat{I}^{(1)}(\hat{\theta}_n)$" for (5), and "$\hat{I}^{(2)}(\hat{\theta}_n)$" for (6); the likelihood-based approach (7) by "W".

As can be seen from Proposition 5.1, already the score functions can be quite complicated. In order to be able to investigate the method $\hat{I}^{(2)}(\hat{\theta}_n)$ based on the observed information, we only consider the Clayton family, for which

$$\begin{aligned}\nabla\nabla^\mathsf{T} l(\boldsymbol{\theta};\boldsymbol{u}) = &-\sum_{k=0}^{d-1}\left(\frac{k}{\theta k+1}\right)^2 + \frac{2}{\theta^2}\left(\frac{t'_\theta(\boldsymbol{u})}{1+t_\theta(\boldsymbol{u})} - \frac{1}{\theta}\log(1+t_\theta(\boldsymbol{u}))\right)\\ &+ (d+1/\theta)\left(\left(\frac{t'_\theta(\boldsymbol{u})}{1+t_\theta(\boldsymbol{u})}\right)^2 - \frac{\sum_{j=1}^d (\log u_j)^2 u_j^{-\theta}}{1+t_\theta(\boldsymbol{u})}\right),\end{aligned}$$

with $t'_\theta(\boldsymbol{u}) = \frac{d}{d\theta} t_\theta(\boldsymbol{u}) = \sum_{j=1}^d (-\log u_j) u_j^{-\theta}$, that is, for which $\hat{I}^{(2)}(\hat{\theta}_n)$ can be easily computed. Our simulation study is based on the sample sizes $n \in \{100, 400\}$ in the dimensions $d \in \{5, 20\}$ for the dependencies $\tau \in \{0.25, 0.5, 0.75\}$. For each of these setups and each of the methods $I(\hat{\theta}_n)$, $\hat{I}^{(1)}(\hat{\theta}_n)$, $\hat{I}^{(2)}(\hat{\theta}_n)$, and W, we determine the proportion of cases among $N = 1000$ replications for which the true parameter is contained in the computed confidence interval. Since the expected information is not known explicitly, we evaluate it by a Monte Carlo simulation based on samples of size 10 000.

Table 3 shows the results of the conducted simulation study. Overall, all methods work comparably well. Note that from a computational point of view, $\hat{I}^{(1)}(\hat{\theta}_n)$ is preferred to $I(\hat{\theta}_n)$ if the latter has to be evaluated based on a Monte Carlo simulation. Furthermore, $\hat{I}^{(2)}(\hat{\theta}_n)$ is typically difficult to evaluate, due to the complicated second order derivatives; the tractable Clayton family is certainly an exception. Even $\hat{I}^{(1)}(\hat{\theta}_n)$ may be (numerically) challenging for some families, as Proposition 5.1 indicates. The likelihood based approach W has several advantages. First, it is typically even simpler to evaluate than $\hat{I}^{(1)}(\hat{\theta}_n)$. Second, it may lead to asymmetric confidence intervals. Finally, by using a re-parameterization, it allows one to construct confidence intervals for quantities such as Kendall's tau or the tail-dependence coefficients (otherwise often obtained from the Delta Method based on the approximate normal distribution).

6 Multi-parameter families

The one-parameter generators of Ali-Mikhail-Haq, Clayton, Frank, Gumbel, and Joe can easily be extended to allow for more parameters, for example, by so-called outer power transformations or even more general generator transformations; see Hofert (2010a), Hofert (2010b), or Hofert (2011a). In this section, we investigate an outer power Clayton copula and the Archimedean GIG family and apply maximum-likelihood estimation for estimating the copula parameters. Both of these families are available via the R

Submitted

Coverage probabilities for Clayton (in %)				Method for obtaining confidence intervals			
$1-\alpha$	n	τ	d	$I(\hat{\theta}_n)$	$\hat{I}^{(1)}(\hat{\theta}_n)$	$\hat{I}^{(2)}(\hat{\theta}_n)$	W
0.95	100	0.25	5	95.6	95.4	96.0	95.6
			20	94.6	95.4	94.9	94.9
		0.5	5	94.1	94.2	94.0	94.0
			20	95.9	96.7	95.8	95.8
		0.75	5	95.7	95.9	95.6	95.7
			20	95.8	95.9	95.9	95.9
	400	0.25	5	94.8	94.9	95.1	95.1
			20	95.7	96.3	96.0	96.0
		0.5	5	95.1	95.2	95.2	95.0
			20	95.6	95.3	95.3	95.2
		0.75	5	94.8	94.5	95.0	94.9
			20	94.7	95.0	94.7	94.7
0.99	100	0.25	5	98.8	98.7	99.1	99.1
			20	98.9	98.4	98.9	98.9
		0.5	5	98.2	98.7	98.4	98.4
			20	99.4	99.2	99.3	99.3
		0.75	5	99.1	99.1	99.2	99.2
			20	98.8	98.7	98.8	98.8
	400	0.25	5	98.7	98.7	98.8	98.8
			20	99.3	99.1	99.3	99.3
		0.5	5	99.0	98.9	99.0	99.0
			20	99.6	99.7	99.6	99.6
		0.75	5	98.7	98.6	98.7	98.7
			20	99.1	99.3	99.1	99.1
0.995	100	0.25	5	99.5	99.5	99.7	99.5
			20	99.4	99.4	99.5	99.5
		0.5	5	99.2	99.2	99.4	99.3
			20	99.8	99.8	99.9	99.9
		0.75	5	99.5	99.5	99.5	99.5
			20	99.6	99.1	99.6	99.5
	400	0.25	5	99.0	99.1	99.3	99.3
			20	99.6	99.4	99.6	99.6
		0.5	5	99.1	99.3	99.2	99.2
			20	99.9	99.8	99.9	99.9
		0.75	5	99.3	99.3	99.3	99.2
			20	99.7	99.6	99.8	99.7

Table 3 Simulated coverage probabilities for Clayton's family based on $N = 1000$ replications.

Submitted

package `nacopula` so that the interested reader can easily follow our calculations. The computations carried out in this section were run on a Mac mini under Mac OS X Version 10.6.6 with a 2.66 GHz Intel Core 2 Duo processor and 4 GB 1067 MHz DDR3 memory. The R version used is 2.12.1.

6.1 Finding initial intervals

Maximizing the log-likelihood l is typically achieved by a numerical routine. These algorithms often require an initial interval (or an initial value, which can be derived from the former). This interval should be sufficiently large in order to contain the optimum, but also sufficiently small in order to find the optimum fast. Furthermore, one should be able to compute an initial interval in a small amount of time in comparison to the actual log-likelihood evaluations required for maximizing the log-likelihood.

For Archimedean families with $\psi_{\boldsymbol{\theta}} \in \Psi_\infty$, the measure of concordance Kendall's tau is a function in $\boldsymbol{\theta}$ which always maps to the unit interval; see, for example, Hofert (2010b, pp. 59). It thus provides an intuitive "distance" in terms of concordance. For one-parameter families, one can thus typically choose an initial interval of the form

$$[\tau^{-1}(\max\{\hat{\tau} - h, \tau_l\}), \tau^{-1}(\min\{\hat{\tau} + h, \tau_u\})],$$

where $h \in [0, 1]$ is suitably chosen with intuitive interpretation as "distance in concordance" and τ_l and τ_u denote lower and upper admissible Kendall's tau for the families considered (in Example 5.2 we used this technique to find an interval on which the log-likelihood is plotted; we took $\hat{\tau}$ as the correct value $\tau = 0.5$, and used $h = 0.01$ and $h = 0.015$ for Clayton's and Gumbel's family, respectively). If the dimension is not too large, one can take the mean of pairwise sample versions of Kendall's tau as estimator $\hat{\tau}$ of Kendall's tau; see Berg (2009), Kojadinovic and Yan (2010), and Savu and Trede (2010) for this estimator. Another option is a multivariate version of Kendall's tau; see Jaworski et al. (2010, pp. 217). A fast way, especially in large dimensions, is to utilize the explicit diagonal maximum-likelihood estimator

$$\hat{\theta}_n^{\mathrm{G}} = \frac{\log d}{\log n - \log(\sum_{i=1}^n - \log Y_i)}, \text{ where } Y_i = \max_{j \in \{1,\dots,d\}} U_{ij},\ i \in \{1, \dots, n\}.$$

for Gumbel's family, see Hofert et al. (2011), and estimate Kendall's tau by $\tau^{\mathrm{G}}(\hat{\theta}_n^{\mathrm{G}})$, where $\tau^{\mathrm{G}}(\theta) = (\theta - 1)/\theta$ denotes Kendall's tau for Gumbel's family as a function in the parameter. Since the optimization for one-parameter families is typically not too time-consuming, one can also just maximize the log-likelihood on a reasonably large, fixed interval, for example $[\tau^{-1}(h_1), \tau^{-1}(h_2)]$, where h_1 and h_2 are suitably chosen constants in the range of τ; see Hofert et al. (2011).

For multi-parameter Archimedean families, the log-likelihood is typically even more challenging to evaluate. An initial interval therefore also serves the purpose of reducing the parameter space to an area where the log-likelihood can be evaluate without numerical problems. The idea we present here to construct initial intervals for multi-parameter families is again based on Kendall's tau. In a first step, we estimate Kendall's tau

Submitted

by $\hat{\tau}_n$. To this end we apply the pairwise Kendall's tau estimator, which, due to the rather complicated log-likelihood evaluations does not take too much run time for the ten-dimensional examples considered below; another option would be to randomly select sub-columns of the data and apply the pairwise Kendall's tau estimator to this sub-data in order to reduce run time. Based on this estimator of Kendall's tau, we then construct an initial rectangle by three points. These points are determined via $\tau^{-1}(\hat{\tau}_n - h_-)$ and $\tau^{-1}(\hat{\tau}_n + h_+)$, that is, via certain positive numbers h_- and h_+ (sufficiently small to ensure that $\hat{\tau}_n - h_-$ and $\hat{\tau}_n + h_+$ are in the range of admissible Kendall's tau). They allow for an intuitive interpretation as "distance in (terms of) concordance" and are independent of the parameterization of the family (since they measure distances in Kendall's tau and not in the underlying copula parameters). Now note that τ^{-1} is not uniquely defined for two- or more-parameter families. It is, however, if one fixes all but one parameter. By starting with one corner of the initial rectangle to be constructed and applying monotonicity properties of τ as a function in its parameters, one can thus construct an initial rectangle around the estimate $\hat{\tau}_n$ of $\tau(\boldsymbol{\theta}_0)$. More details are given in Sections 6.2 and 6.3 for the two-parameter Archimedean families investigated.

6.2 Outer power copulas

If $\psi \in \Psi_\infty$, so is $\tilde{\psi}(t) = \psi(t^{1/\beta})$ for all $\beta \in [1, \infty)$, since the composition of a completely monotone function with a non-negative function that has a completely monotone derivative is again completely monotone; see Feller (1971, p. 441). The copula family generated by $\tilde{\psi}$ is referred to as *outer power family*.

The generator derivatives of $\tilde{\psi}(t) = \psi(t^{1/\beta})$ can be accessed with a formula about derivatives of compositions which dates back at least to Schlömilch (1846). According to this formula,

$$(-1)^d \tilde{\psi}^{(d)}(t) = P^{\text{op}}(t^{1/\beta})/t^d,\ d \in \mathbb{N}, \quad \text{where} \quad P^{\text{op}}(x) = \sum_{k=1}^{d} a_{dk}^{\text{G}}(\beta)(-1)^k \psi^{(k)}(x) x^k.$$

Via (2) and the form of a_{dk}^{G} given in Theorem 3.2 (4) one can thus easily derive the density of an outer power copula.

For sampling $\tilde{V} \sim \tilde{F} = \mathcal{LS}^{-1}[\tilde{\psi}]$, Hofert (2011a) derived the stochastic representation

$$\tilde{V} = SV^\beta,\ S \sim \text{S}(1/\beta, 1, \cos^\beta(\pi/(2\beta)), \mathbb{1}_{\{\beta=1\}}; 1),\ V \sim F = \mathcal{LS}^{-1}[\psi].$$

Note that $\tilde{V}$ can easily be sampled via the R package `nacopula` for all ψ given in Table 1.

We consider the case where ψ is Clayton's generator, so we obtain the two-parameter *outer power Clayton copula* with generator $\tilde{\psi}(t) = (1 + t^{1/\beta})^{-1/\theta}$. This copula, which generalizes the Clayton family, was successfully applied in Hofert and Scherer (2011) in the context of pricing collateralized debt obligations. For this copula, Kendall's tau and the tail-dependence coefficients are given explicitly by

$$\tau = \tau(\theta, \beta) = 1 - \frac{2}{\beta(\theta + 2)}, \quad \lambda_L = 2^{-1/(\beta\theta)}, \quad \lambda_U = 2 - 2^{1/\beta}. \tag{8}$$

Submitted

Note the possibility to have upper tail dependence for this copula, which is not possible for a Clayton copula.

The following algorithm describes a procedure for finding an initial interval for outer power Clayton copulas. The algorithm can easily be adapted to other outer power copulas, given that the base family (the family generated by ψ) is positively ordered in its parameter and admits a sufficiently large range of Kendall's tau.

Algorithm 6.1

(1) Choose $h_-, h_+ \geq 0$, and $\varepsilon > 0$.

(2) Let the smallest β be denoted by $\beta_l = 1$.

(3) Solve $\tau(\theta_u, \beta_l) = \min\{\hat{\tau}_n + h_+, 1 - \varepsilon\}$ with respect to θ_u.

(4) Solve $\tau(\theta_l, \beta_l) = \max\{\hat{\tau}_n - h_-, \varepsilon\}$ with respect to θ_l.

(5) Solve $\tau(\theta_l, \beta_u) = \min\{\hat{\tau}_n + h_+, 1 - \varepsilon\}$ with respect to β_u.

(6) Return the initial interval $I = [(\theta_l, \beta_l)^\mathsf{T}, (\theta_u, \beta_u)^\mathsf{T}]$.

The idea behind Algorithm 6.1 is to construct an initial rectangle by three points. First, the lower-right endpoint of the rectangle is constructed. Since $\tau(\theta, \beta)$ is an increasing function in both θ and β, the largest θ and the smallest β, that is, $(\theta_u, \beta_l)^\mathsf{T}$, are chosen such that Kendall's tau equals $\hat{\tau}_n$ plus a small "distance in concordance" $h_+ \geq 0$ to ensure that θ_u is indeed an upper bound for θ. The truncation done by $\varepsilon > 0$ is to obtain an admissible Kendall's tau range. Second, the lower-left endpoint is found. The monotonicity of τ justifies determining the minimal value θ_l for θ such that $\tau(\theta_l, \beta_l) = \max\{\hat{\tau}_n - h_-, \varepsilon\}$, where $h_- \geq 0$ is suitably chosen, similar to h_+. In the third and final step, the upper-left endpoint of the initial rectangle is determined. The maximal value β_u for β is determined in a similar fashion to the first step. Note that all equations can be solved explicitly due to the explicit form of Kendall's tau as given in (8).

To access the performance of the maximum-likelihood estimator, we generate $N = 1000$ times $n = 100$ realizations of i.i.d. random vectors following d-dimensional outer power Clayton copulas. For demonstration purposes, we consider $d = 10$. Furthermore, we consider three setups of dependencies: $\boldsymbol{\theta} = (\theta, \beta)^\mathsf{T} = (1/3, 8/7)^\mathsf{T}$ resulting in a Kendall's tau of 0.25; $\boldsymbol{\theta} = (1, 4/3)^\mathsf{T}$ with corresponding Kendall's tau equal to 0.5; and $\boldsymbol{\theta} = (2, 2)^\mathsf{T}$ with Kendall's tau equal to 0.75. For finding initial intervals, Algorithm 6.1 is applied with $\varepsilon = 0.005$, $h_- = 0.4$, and $h_+ = 0$. The results are summarized in Table 4, where "RMSE" denotes the root mean squared error as before and "MUT" denotes the mean user time (in seconds).

Figure 4 shows a wire-frame plot (left) of the negative log-likelihood of a sample of size $n = 100$ for the setup $\boldsymbol{\theta} = (1, 4/3)^\mathsf{T}$ ($\tau = 0.5$) and the corresponding level plot (right). Both plots have the initial interval determined by Algorithm 6.1 as domain and show both the true value $\boldsymbol{\theta}_0 = (\theta_0, \beta_0)^\mathsf{T}$ and the optimum $\hat{\boldsymbol{\theta}}_n = (\hat{\theta}_n, \hat{\beta}_n)^\mathsf{T}$ as determined by the optimizer. Figure 5 shows profile likelihood plots for the two parameters θ and β.

Submitted

			$\hat{\theta}_n$			$\hat{\beta}_n$			
n	τ	θ	Bias	RMSE	β	Bias	RMSE	#	MUT
100	0.25	0.3333	0.0073	0.0609	1.1429	−0.0017	0.0429	42	0.7 s
100	0.5	1.0000	0.0082	0.1050	1.3333	−0.0003	0.0613	39	0.6 s
100	0.75	2.0000	0.0107	0.1786	2.0000	−0.0012	0.1088	47	0.7 s
500	0.25	0.3333	0.0025	0.0276	1.1429	−0.0014	0.0189	42	1.2 s
500	0.5	1.0000	0.0026	0.0451	1.3333	−0.0013	0.0268	38	1.1 s
500	0.75	2.0000	0.0031	0.0753	2.0000	−0.0013	0.0483	49	1.3 s

Table 4 Summary statistics for estimating two-parameter outer power Clayton copulas.

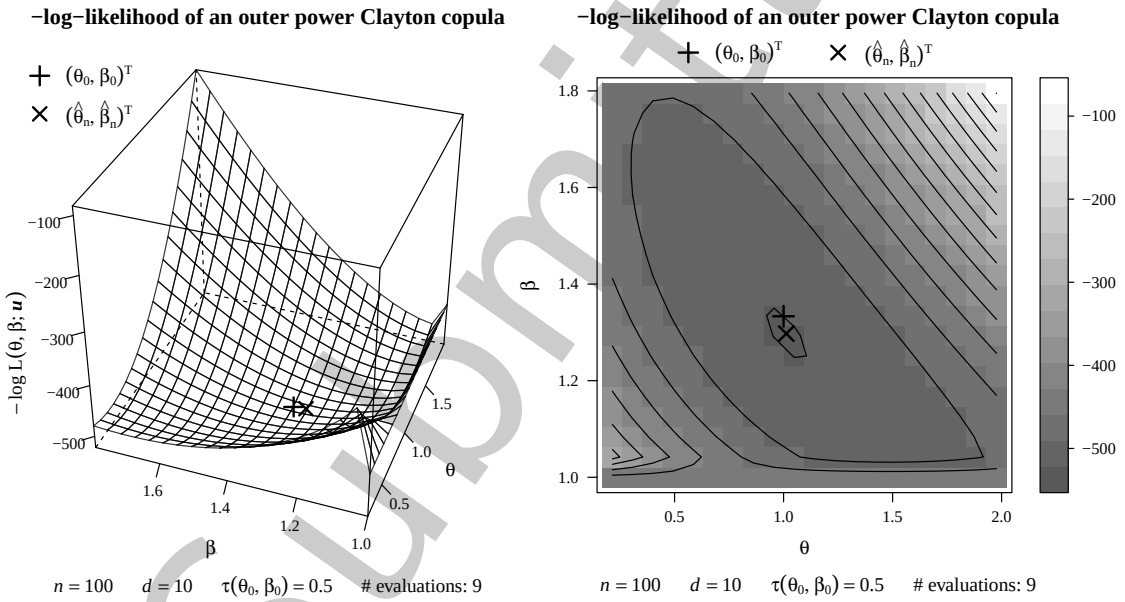

Figure 4 A wire-frame plot (left) and corresponding level plot (right) of the negative log-likelihood function for an outer power Clayton copula for a sample of size $n = 100$ for $\boldsymbol{\theta} = (1, 4/3)^\mathsf{T}$ ($\tau = 0.5$) with the computed initial interval as domain.

Submitted

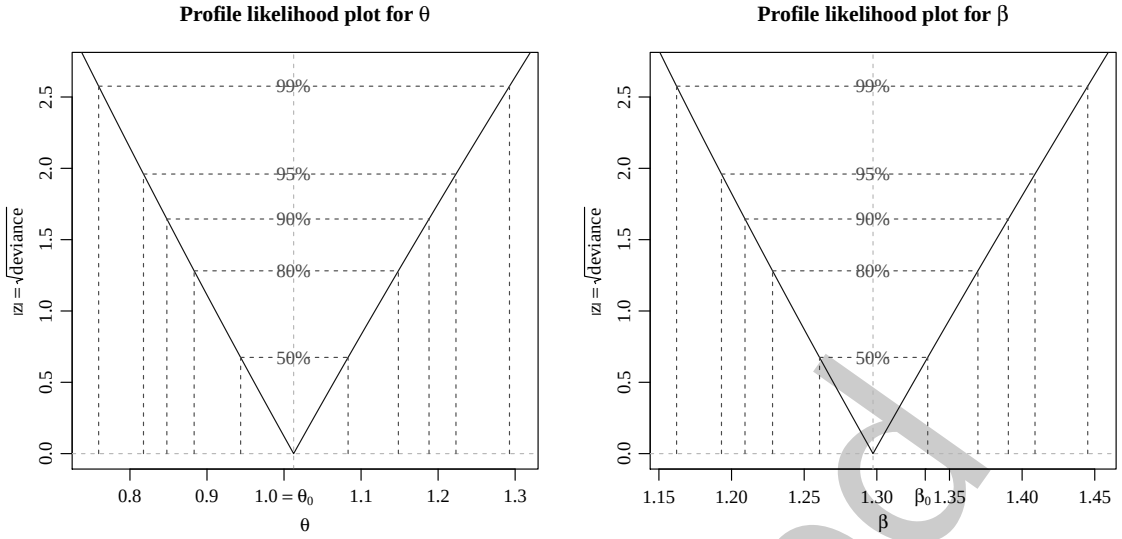

Figure 5 Profile likelihood plot for θ (left) and β (right).

6.3 The GIG family

An Archimedean family which naturally allows for two parameters can be constructed as follows. We start with the density of a *generalized inverse Gaussian distribution* $\mathrm{GIG}(\nu, \phi, \chi)$, given by

$$g(x; \nu, \phi, \chi) = \frac{(\phi/\chi)^{\nu/2}}{2K_\nu(\sqrt{\phi\chi})} x^{\nu-1} \exp(-(\chi/x + \phi x)/2),\ x \in (0, \infty).$$

Here, $\nu \in \mathbb{R}$ with: $\phi \in [0, \infty)$, $\chi \in (0, \infty)$, if $\nu \in (-\infty, 0)$; $\phi, \chi \in (0, \infty)$, if $\nu = 0$; and $\phi \in (0, \infty)$, $\chi \in [0, \infty)$, if $\nu \in (0, \infty)$; see McNeil et al. (2005, p. 497). The function $K_\nu(t) = \int_0^\infty \cosh(\nu x) \exp(-t \cosh(x))\, dx$ denotes the *modified Bessel function of the third kind* with parameter ν. It is decreasing in t and symmetric about zero in ν. Furthermore, it is increasing in ν if $\nu \in (0, \infty)$. Another important property is

$$\lim_{t \searrow 0} t^\nu K_\nu(t) = 2^{\nu-1}\Gamma(\nu) \text{ if } \nu \in (0, \infty). \tag{9}$$

Note that for a $\psi \in \Psi_\infty$, the generator $\psi(ct)$ generates the same Archimedean copula as $\psi(t)$ for all $c \in (0, \infty)$. Letting $c = \phi/2$ and $\theta = \sqrt{\phi\chi}$ leads to a comparably simple form of the generator of an *Archimedean GIG copula* with parameter vector $\boldsymbol{\theta} = (\nu, \theta)^\mathsf{T}$, given by

$$\psi(t) = (1 + t)^{-\nu/2} K_\nu(\theta\sqrt{1 + t})/K_\nu(\theta),\ t \in [0, \infty),\ \nu \in \mathbb{R},\ \theta \in (0, \infty). \tag{10}$$

If we let

$$h_{\nu_1,\nu_2,\theta}(t) = \frac{(\theta\sqrt{1+t})^{\nu_1} K_{\nu_1}(\theta\sqrt{1+t})}{\theta^{\nu_2} K_{\nu_2}(\theta)},\ \nu_1, \nu_2 \in \mathbb{R},\ \theta \in (0, \infty),\ t \in [0, \infty),$$

Submitted

one obtains from (9) that $\lim_{\theta \searrow 0} h_{\nu_1,\nu_2,\theta}(t) = 2^{\nu_1-\nu_2}\Gamma(\nu_1)/\Gamma(\nu_2)$ for every $\nu_1, \nu_2 \in (0,\infty)$ and $t \in [0,\infty)$. Since ψ can be written as $\psi(t) = (1+t)^{-\nu} h_{\nu,\nu,\theta}(t)$, one obtains as limiting case a $\Gamma(\nu,1)$ distribution for $\theta \searrow 0$ if $\nu \in (0,\infty)$, that is, a Clayton copula with parameter $1/\nu$.

The density f of $F = \mathcal{LS}^{-1}[\psi]$ is given by

$$f(x;\nu,\theta) = \frac{(2x)^{\nu-1}}{\theta^\nu K_\nu(\theta)} \exp(-(\theta^2/(2x) + 2x)/2),\ x \in (0,\infty),$$

so that $V = X/2 \sim F$ with $X \sim \mathrm{GIG}(\nu, 1, \theta^2)$. The GIG distribution can easily be sampled with the R package `Runuran`. For numerically computing ψ^{-1}, note that $K_\nu(y)/K_\nu(x) < \exp(x-y)(y/x)^\nu$ for all $\nu \in (-1/2,\infty)$ and $0 < x < y < \infty$ (see Paris (1984)), so that $[0, (1-\log(t)/\theta)^2 - 1]$ is an initial interval for searching $\psi^{-1}(t)$ for all $\nu \in (-1/2,\infty)$.

Kendall's tau and the coefficients of tail dependence of a GIG copula are given by

$$\tau = \tau(\nu,\theta) = 1 - \int_0^\infty t(\theta h_{\nu+1,\theta}(t)/K_\nu(\theta))^2\, dt, \quad \lambda_L = \lambda_U = 0,$$

where $h_{\nu+k,\theta}(t) = K_{\nu+k}(\theta\sqrt{1+t})/\sqrt{1+t}^{\nu+k}$, $t \in [0,\infty)$, $\theta \in (0,\infty)$, and $\nu, k \in \mathbb{R}$. For computing the tail dependence coefficients, consider $\psi'(2t)/\psi'(t)$ for the limits $t \downarrow 0$ and $t \uparrow \infty$, and use $K_\nu(t) \approx \sqrt{\pi/(2t)}\exp(-t)$ (valid for $t \gg |\nu^2 - 1/4|$) in the latter case. A numerically stable evaluation of the integral formula for Kendall's tau for small $\nu > 0$ based on the Clayton limit is given in the R package `nacopula`. Note that as numerical results indicate, Kendall's tau is decreasing in both ν and θ if $\nu \in [0,\infty)$; see Figure 6 (left). If $\nu \in (0,\infty)$, the limit for Kendall's tau as $\theta \searrow 0$ is $1/(1+2\nu)$ which equals Kendall's tau for Clayton's family with parameter $1/\nu$. Figure 6 (right) shows a scatter plot of 1000 bivariate vectors of random variates following a GIG family with $\boldsymbol{\theta} = (0.05, 0.0968)^\mathsf{T}$ with corresponding Kendall's tau equal to 0.5.

One advantage of the GIG family is that the generator derivatives take on a comparably simple form, which can be represented as

$$(-1)^d \psi^{(d)}(t) = \frac{h_{\nu+d,\nu,\theta}(t)}{2^d(1+t)^{\nu+d}},\ t \in (0,\infty),\ d \in \mathbb{N}_0.$$

This can be easily derived by differentiating ψ under the integral sign and interpreting the resulting integrand as the density of a $\mathrm{GIG}(\nu+d, 2(1+t), \theta^2/2)$ distribution which integrates to one. Via (2), one then easily finds the form of the log-density of a GIG family, given by

$$\begin{aligned}
\log c(\boldsymbol{u}) &= \log h_{\nu+d,\theta}(t_{\boldsymbol{\theta}}(\boldsymbol{u})) + (d-1)\log K_\nu(\theta) - \sum_{j=1}^d \log h_{\nu+1,\theta}(\psi^{-1}(u_j)), \\
&= -(\nu+d)\log(1+t_{\boldsymbol{\theta}}(\boldsymbol{u})) + \log h_{\nu+d,\nu,\theta}(t_{\boldsymbol{\theta}}(\boldsymbol{u})) + (\nu+1)\sum_{j=1}^d \log(1+\psi^{-1}(u_j)) \\
&\quad - \sum_{j=1}^d \log h_{\nu+1,\nu,\theta}(\psi^{-1}(u_j)).
\end{aligned}$$

Submitted

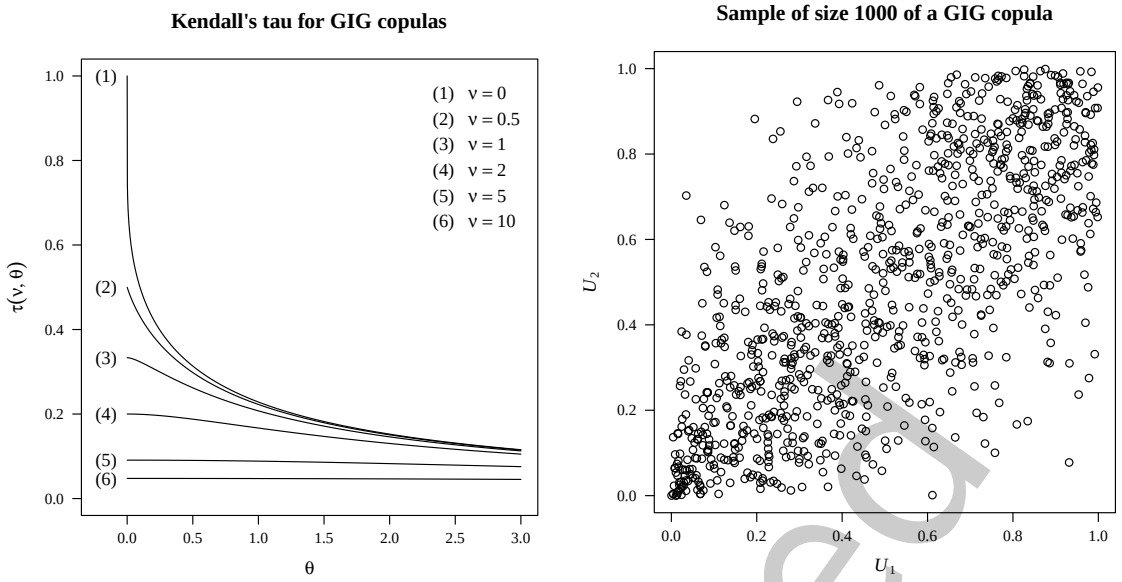

Figure 6 Kendall's tau for the GIG family as a function in θ for given ν (left). A sample of 1000 bivariate vectors of random variates following a GIG family with $\boldsymbol{\theta} = (0.05, 0.0968)^\mathsf{T}$ and corresponding Kendall's tau equal to 0.5 (right).

The following algorithm describes a procedure for finding an initial interval for GIG copulas with $\nu \in [0, \infty)$. The idea behind this algorithm is similarly to those of Algorithm 6.1. However, it takes into account that $\tau(\nu, \theta)$ is decreasing in both parameters $\nu \in [0, \infty)$ and $\theta \in (0, \infty)$ and thus first determines the upper-left, then the lower-left, and finally the lower-right endpoint of the initial rectangle. As before, $\hat{\tau}_n$ denotes an estimator of Kendall's tau, taken as the pairwise Kendall's tau estimator. Note that for $\nu = \nu_l = 0$, the range of Kendall's tau as a function in θ is $(0, 1]$; see Figure 6 (left). Furthermore, for $\theta = \theta_l$ sufficiently small, the range of Kendall's tau as a function in ν is $(0, 1)$.

Algorithm 6.2

(1) Choose $h_-, h_+ \geq 0$, and $\varepsilon > 0$.
(2) Let the smallest ν be denoted by $\nu_l = 0$.
(3) Solve $\tau(\nu_l, \theta_u) = \max\{\hat{\tau}_n - h_-, \varepsilon\}$ with respect to θ_u.
(4) Solve $\tau(\nu_l, \theta_l) = \min\{\hat{\tau}_n + h_+, 1 - \varepsilon\}$ with respect to θ_l.
(5) Solve $\tau(\nu_u, \theta_l) = \max\{\hat{\tau}_n - h_-, \varepsilon\}$ with respect to ν_u.
(6) Return the initial interval $I = [(\nu_l, \theta_l)^\mathsf{T}, (\nu_u, \theta_u)^\mathsf{T}]$.

We generate $N = 200$ times $n = 100$ realizations of i.i.d. random vectors in $d = 10$ dimensions following GIG copulas with parameters $\boldsymbol{\theta} = (\nu, \theta)^\mathsf{T} = (0.1, 0.8333)^\mathsf{T}$ resulting in a Kendall's tau of 0.25, $\boldsymbol{\theta} = (0.05, 0.0968)^\mathsf{T}$ with corresponding Kendall's tau equal to 0.5, and $\boldsymbol{\theta} = (0.01, 0.0012)^\mathsf{T}$ with Kendall's tau equal to 0.75 (the choice of N is made

Submitted

solely to the larger run time for this family). For finding initial intervals, Algorithm 6.2 is applied with $\varepsilon = 0.005$ and $h_- = h_+ = 0.15$. The results are summarized in Table 5. Note that especially under weak concordance, the GIG family requires a larger sample size n in order for the bias to be small.

			$\hat{\nu}_n$			$\hat{\theta}_n$			
n	τ	ν	Bias	RMSE	θ	Bias	RMSE	#	MUT
100	0.25	0.1000	0.1622	0.3769	0.8333	−0.0178	0.1395	57	12.3 s
100	0.5	0.0500	0.0065	0.0538	0.0968	0.0002	0.0134	55	12.7 s
100	0.75	0.0100	0.0198	0.0541	0.0012	−0.0001	0.0005	93	39.3 s
500	0.25	0.1000	0.0480	0.1680	0.8333	−0.0003	0.0562	63	67.0 s
500	0.5	0.0500	0.0034	0.0295	0.0968	0.0003	0.0063	53	60.4 s
500	0.75	0.0100	0.0057	0.0286	0.0012	−0.0000	0.0003	102	168.3 s

Table 5 Summary statistics for estimating two-parameter GIG copulas.

7 Conclusion

We presented explicit functional forms for the generator derivatives of well-known Archimedean copulas. These explicit formulas are of interest for several reasons. Apart from being able to express various important quantities such as conditional distributions or the Kendall distribution function explicitly, the generator derivatives allow us to apply maximum-likelihood estimation for estimating the parameter vectors of various Archimedean copulas, even in large dimensions such as $d = 100$. The excellent performance in terms of both precision and run time of maximum likelihood estimation was shown in Hofert et al. (2011) in a large-scale comparison with various other estimators up to dimension $d = 100$. In the present work, we presented the theoretical details and showed that maximum-likelihood estimation is also feasible for multi-parameter Archimedean families. Furthermore, we showed that the mean squared error MSE is decreasing in the dimension d and that this decrease is of the same order as the decrease in the sample size n, that is, MSE $\propto 1/(nd)$. We also constructed initial intervals for the likelihood optimization. Moreover, we obtained likelihood-based confidence intervals for the parameter vector and compared them to information-based confidence intervals for the Clayton family where the Fisher information is comparably easy to compute. A transparent implementation of the presented results is given in the open source R package `nacopula`, so that the interested reader can easily follow our calculations.

Acknowledgements

The authors would like to thank Khristo Boyadzhiev (Ohio Northern University) for introducing us to and guiding us through the fascinating world of exponential polynomials.

Submitted

References

P. Barbe, C. Genest, K. Ghoudi, and B. Rémillard. On Kendall's Process. *Journal of Multivariate Analysis*, 58:197–229, 1996.

D. Berg. Copula goodness-of-fit testing: an overview and power comparison. *The European Journal of Finance*, 2009. URL http://www.informaworld.com/10.1080/13518470802697428.

D. Berg and K. Aas. Models for construction of multivariate dependence – A comparison study. *The European Journal of Finance*, 15(7):639–659, 2009.

P. J. Bickel and K. A. Doksum. *Copula Theory and Its Applications*. Prentice Hall, 2 edition, 2000.

K. N. Boyadzhiev. Exponential Polynomials, Stirling Numbers, and Evaluation of Some Gamma Integrals. *Abstract and Applied Analysis*, 2009, 2009.

D. R. Cox and D. V. Hinkley. *Theoretical Statistics*. Chapman and Hall, 1974.

A. C. Davison. *Statistical Models*. Cambridge Series in Statistical and Probabilistic Mathematics, 2003.

B. Efron and D. V. Hinkley. Assessing the accuracy of the maximum likelihood estimator: Observed versus expected Fisher information. *Biometrika*, 65(3):457–487, 1978.

P. Embrechts and M. Hofert. On Archimedean copulas and a non-parametric estimation method. *TEST*, 2011. doi: 10.1007/s11749-011-0252-4. in press.

W. Feller. *An Introduction to Probability Theory and Its Applications*, volume 2. Wiley, 2 edition, 1971.

C. Genest, K. Ghoudi, and L.-P. Rivest. A semiparametric estimation procedure of dependence parameters in multivariate families of distributions. *Biometrika*, 82(3): 543–552, 1995.

C. Genest, J. F. Quessy, and B. Rémillard. Goodness-of-fit procedures for copula models based on the probability integral transformation. *Scandinavian Journal of Statistics*, 33:337–366, 2006.

C. Genest, B. Rémillard, and D. Beaudoin. Goodness-of-fit tests for copulas: A review and a power study. *Insurance: Mathematics and Economics*, 44:199–213, 2009.

C. Hering and M. Hofert. Goodness-of-fit tests for Archimedean copulas in large dimensions. 2011. submitted.

M. Hofert. Construction and sampling of nested Archimedean copulas. In F. Durante, W. Härdle, P. Jaworski, and T. Rychlik, editors, *Copula Theory and Its Applications, Proceedings of the Workshop held in Warsaw 25-26 September 2009*, pages 147–160. Springer, 2010a. doi: 10.1007/978-3-642-12465-5_7.

M. Hofert. *Sampling Nested Archimedean Copulas with Applications to CDO Pricing*. Südwestdeutscher Verlag für Hochschulschriften AG & Co. KG, 2010b. ISBN 978-3-8381-1656-3. PhD thesis.

M. Hofert. Efficiently sampling nested Archimedean copulas. *Computational Statistics & Data Analysis*, 55:57–70, 2011a. doi: 10.1016/j.csda.2010.04.025.

M. Hofert. A stochastic representation and sampling algorithm for nested Archimedean copulas. *Journal of Statistical Computation and Simulation*, 2011b. doi: 10.1080/00949655.2011.574632. in press.

Submitted

M. Hofert and M. Scherer. CDO pricing with nested Archimedean copulas. *Quantitative Finance*, 11(5):775–787, 2011. doi: 10.1080/14697680903508479.

M. Hofert and F. Vrins. Sibuya copulas. 2011. in progress, early version: http://arxiv.org/pdf/1008.2292.

M. Hofert, M. Mächler, and A. J. McNeil. Estimators for Archimedean copulas in high dimensions: A comparison. 2011. in progress.

P. Jaworski, F. Durante, W. K. Härdle, and T. Rychlik, editors. *Copula Theory and Its Applications*, volume 198 of *Lecture Notes in Statistics – Proceedings*. Springer, 2010.

H. Joe. *Multivariate Models and Dependence Concepts*. Chapman & Hall/CRC, 1997.

H. Joe and T. Hu. Multivariate Distributions from Mixtures of Max-Infinitely Divisible Distributions. *Journal of Multivariate Analysis*, 57:240–265, 1996.

H. Joe and J. J. Xu. The Estimation Method of Inference Functions for Margins for Multivariate Models. *Technical Report no. 166, Department of Statistics, University of British Columbia*, 1996.

G. Kim, M. J. Silvapulle, and P. Silvapulle. Comparison of semiparametric and parametric methods for estimating copulas. *Computational Statistics & Data Analysis*, 51:2836–2850, 2007.

I. Kojadinovic and J. Yan. Modeling Multivariate Distributions with Continuous Margins Using the copula R Package. *Journal of Statistical Software*, 34(9):1–20, 2010.

E. L. Lehmann and G. Casella. *Theory of Point Estimation*. Springer, 2 edition, 1998.

A. J. McNeil and J. Nešlehová. Multivariate Archimedean copulas, d-monotone functions and l_1-norm symmetric distributions. *The Annals of Statistics*, 37(5b):3059–3097, 2009.

A. J. McNeil, R. Frey, and P. Embrechts. *Quantitative Risk Management: Concepts, Techniques, Tools*. Princeton University Press, 2005.

R. B. Nelsen. *An Introduction to Copulas*. Springer, 2007.

W. K. Newey and D. McFadden. Large sample estimation and hypothesis testing. In R. F. Engle and D. L. McFadden, editors, *Handbook of Econometrics*, pages 2111–2245. Elsevier North Holland, 1994.

R. B. Paris. An inequality for the Bessel function $J_\nu(\nu x)$. *SIAM Journal on Mathematical Analysis*, 15(1):203–205, 1984.

V. K. Rohatgi. *An introduction to probability theory and mathematical statistics*. Wiley, 1976.

C. Savu and M. Trede. Hierarchies of Archimedean copulas. *Quantitative Finance*, 10(3): 295–304, 2010.

M. J. Schervish. *Theory of Statistics*. Springer, 1995.

O. Schlömilch. Allgemeine Sätze für eine Theorie der höheren Differential-Quotienten. *Archiv der Mathematik und Physik*, 7:204–214, 1846.

R. J. Serfling. *Approximation Theorems Of Mathematical Statistics* . Wiley-Interscience, 1980.

D. Shi. Fisher information for a multivariate extreme value distribution. *Biometrika*, 82 (3):644–649, 1995.

A. W. van der Vaart. *Asymptotic Statistics*. Cambridge University Press, 2000.

G. N. F. Weiß. Copula parameter estimation: numerical considerations and implications for risk management. *The Journal of Risk*, 13(1):17–53, 2010.

Submitted

F. Wu, E. A. Valdez, and M. Sherris. Simulating Exchangeable Multivariate Archimedean Copulas and its Applications. *Communications in Statistics – Simulation and Computation*, 36(5):1019–1034, 2007.

Submitted